arXiv:1702.03959v1 [cs.DS] 13 Feb 2017

How large is your graph?

Varun Kanade*[1,2], Frederik Mallmann-Trenn†[3,4], and Victor Verdugo‡[3,5]

[1]Department of Computer Science, University of Oxford
[2]The Alan Turing Institute
[3]Départment d'Informatique, École normale supérieure
[4]Department of Computer Science, Simon Fraser University
[5]Departmento de Ingeniería Industrial, Universidad de Chile

February 15, 2017

Abstract

We consider the problem of estimating the graph size, where one is given only local access to the graph. We formally define a query model in which one starts with a *seed* node and is allowed to make queries about neighbours of nodes that have already been seen. In the case of undirected graphs, an estimator of Katzir *et al.* (2014) based on a sample from the stationary distribution π uses $O\left(\frac{1}{\|\pi\|_2}+d_{\text{avg}}\right)$ queries; we prove that this is tight. In addition, we establish this as a lower bound even when the algorithm is allowed to crawl the graph arbitrarily; the results of Katzir *et al.* give an upper bound that is worse by a multiplicative factor $t_{\text{mix}}\cdot\log(n)$.

The picture becomes significantly different in the case of directed graphs. We show that without strong assumptions on the graph structure, the number of nodes cannot be predicted to within a constant multiplicative factor without using a number of queries that are at least linear in the number of nodes; in particular, rapid mixing and small diameter, properties that most real-world networks exhibit, do not suffice. The question of interest is whether any algorithm can beat breadth-first search. We introduce a new parameter, generalising the well-studied conductance, such that if a suitable bound on it exists and is known to the algorithm, the number of queries required is sublinear in the number of edges; we show that this is tight.

*varunk@cs.ox.ac.uk
†frederik.mallmann-trenn@ens.fr
‡vverdugo@dii.uchile.cl

1 Introduction

Networks contain a wealth of information and studying properties of networks may yield important insights. However, most networks of interest are very large and ordinary users may have rather restricted access to them. One of most basic questions about networks is the number of nodes contained in them. For example, the number of pages on the world wide web (WWW) is estimated to be just shy of 50 billion at the time of writing.[1] Facebook currently reports having about one and three quarter billion users;[2] Twitter reports having about 300 million active users. It is undesirable to rely on a small number of sources for such information. At times we might be interested in more specific graphs for which there is no public information available at all. Is there a way to estimate the total number of nodes using rather limited access to these graphs?

Our model is motivated by the kind of access ordinary users may have to graphs of interest. For example, most social network companies provide some sort of an application programming interface (API). In the case of the world wide web, one option is to simply crawl. The graph query access models we use are formally defined in Section 2.1. For now, we mention four specific networks, each of which captures an important modelling aspect. First, Facebook is an undirected graph of friendships, and as long as privacy settings allow it, it is possible to request the number of friends for a given user and the identity of the friends. Second, the world wide web, which is a directed graph—it is possible to extract out-links on a given webpage; however, there is no obvious method to access all in-links. The third is what we refer to as the "fan" network—for a specific user it is possible to query who she is a fan of, however in terms of her fans only the number is revealed.[3] And finally, the fourth is the twitter network, which is directed, however both the followers and followees of a given user are accessible, as far as privacy settings allow. Obviously, the method to estimate the number of nodes may be rather different in each case. It is worth mentioning that the twitter network can essentially be treated as an undirected graph, however it still leaves open the possibility that differentiating in-links and out-links leads to better estimators.

While memory and computational requirements do act as constraints when dealing with large graphs, possibly the most important one is the rate limits set on queries made using the API. Even when crawling the web, there is the risk of simply being blocked if large volumes of requests are sent to the same server. Thus, the most scarce resource in this instance is the number of queries made about the graph. Our query model counts these costs strictly—essentially every time a new node is discovered, its degree is revealed at unit cost and a unit cost is incurred for every neighbour requested.

1.1 Related Work

There is a large body of literature on estimating statistical properties of graphs, reflecting the relevance of and interest in studying complex networks. Much of this work is in the more applied literature and in particular, we were unable to pin down a precisely defined query model. The work most closely related to ours is that of Katzir *et al.* [12]. They consider a model where a random sample drawn from the stationary distribution of the random walk is available. If the random walk mixes rapidly and a suitable bound on the mixing time is known, this can be simulated in the models we consider. They show that $O\left(\frac{1}{\|\pi\|_2}+d_{\text{avg}}\right)$ queries suffice, where π is the stationary distribution and d_{avg} is the average degree. We show that this is tight when given access to a sample from the stationary distribution. When only neighbour queries are allowed, we show that this bound is tight up to a multiplicative factor of the mixing

[1] www.worldwidewebsize.com

[2] Here, billion refers to one thousand million, not a million million.

[3] Although it is not obvious which of the extant networks has this property, there are close approximations such as Blogger. In any case, it is a very natural intermediate model between the second and the fourth.

time and other polylogarithmic factors. It is worth mentioning what this bound actually yields in graphs where the average degree is small, something common to most real world networks. It is always the case that $\|\pi\|_2^{-1} \leq \sqrt{n}$, where equality holds in the case of regular graphs. Thus, the number of queries required is significantly sublinear in the number of nodes. For graphs with power law degree distributions with parameter $\beta = 2$, Katzir *et al.* calculated that $\|\pi\|_2^{-1} = O(n^{1/4} \log n)$. See Section 3.1 for further discussion.

In more recent work, Hardiman and Katzir give another estimator based on counting shared neighbours [11]. They give a slightly better bound on the number of nodes sampled than the earlier work of Katzir *et al.* [12]. However, it is unclear that this can be implemented efficiently in the query model in our paper. Cooper *et al.* [4] obtain estimators for the number of edges, triangles and nodes. The estimators are based on random walks over the graph, but in particular, for estimating the number of nodes, the transition probabilities are not inversely proportional to the degree. Their estimator relies on counting the time required for the k^{th} return to a particular vertex. It seems unlikely that either their random walk or their estimator can be implemented in a query efficient manner. Musco *et al.* [17] develop a distributed collision based approach where several random walks traverse the graph and counts the number of collisions with other random walks. Dasgupta *et al.* [6] provide an estimator of the average degree that uses $O(\log U \log\log U)$ samples, where U is an upper bound on the maximum degree. Somewhat surprisingly, this seems to be an easier task than estimating the number of nodes.

Cooper and Frieze [3] estimate the graph size in a dynamic setting where the graph grows over time and there exists an agent that is allowed to move through the network every time a new node arrives. They prove that this agent visits asymptotically a constant fraction of the vertices. If one has access to the complete neighbourhood of a node when it is visited, it is shown by Mihail *et al.* that the expected time in which a walk discovers a power law random graph is sublinear [16]. There has been some older theoretical work on estimating the size of graphs motivated primarily by search algorithms [13, 15]. However the model of graph access is unrelated and they do not present bounds on the estimates.

Related to the question of estimating graph properties is that of testing whether graphs have specific properties. Property testing has received much attention in the last two decades with properties of undirected graphs being one of the important areas of focus (see *e.g.,* [9]). Directed graphs have received somewhat less attention; there are two main query models, *unidirectional* where only out-neighbours may be queried, and *bidirectional* where both in and out-neighbours may be queried [1]. Bender and Ron showed that there exist properties such as strong connectivity, where the query complexity may be either $O(1)$ or $\Omega(\sqrt{n})$ depending on the model used, even when allowing two-sided error. More recently, Czumaj *et al.* have shown that if a property can be tested with constant query complexity in the bidirectional setting, it can be tested with sublinear query complexity in the unidirectional model [5]. Although, these query models are closely related to the ones in this paper, a crucial aspect exploited by most property testing algorithms is the ability to sample nodes uniformly at random, something that is not available in our setting.

A closely related line of work is that on estimating properties of distributions from samples; of most interest, in our case is the support size. This problem has a long history going back at least to Good and Turing [10]. We only mention a few recent relevant results here. Given access to uniformly sampled elements from a set, $O(\sqrt{n})$ samples suffice to derive a good estimate of the set size using the birthday problem [8]. However, it is not clear how to sample from the uniform distribution over the nodes of the graph in the query model we consider. Valiant and Valiant show that support size can be estimated to a good accuracy using $O(n/\log n)$ samples for any distribution [18, 19]. However, their result requires that any element in the support has $\Omega(1/n)$ probability mass, something that is not true for the stationary distribution of a random walk on a graph. (In the case of directed graphs this problem becomes even more severe, since some nodes may have exponentially small stationary probability mass.)

1.2 Our contributions

Firstly, our contribution is to express the problem of estimating graph properties formally. As discussed in the introduction, networks of interest vary significantly in terms of what access might be easily available to an ordinary user. Keeping in mind the examples of Facebook, the web, the fan-network and Twitter, we introduce different types of oracles that provide access to the graph.

The focus of this work is on estimating the number of nodes. For any $\varepsilon > 0$ and $\delta > 0$, we say that an algorithm (with access to a query oracle) provides an ε-accurate estimate of the number of nodes, if with probability at least $1-\delta$ it outputs $\widehat{n}$, such that $|n-\widehat{n}| \leq \varepsilon n$. The main quantity to be optimised is the number of oracle queries, though all algorithms considered in this paper are also computationally highly efficient. Allowed queries are defined precisely in Section 2.1. Here, we point out that a unit cost must be paid for every disclosed neighbour as well as to know the degree of a node. All algorithms have access to the *identifier* of one *seed* node in the graph to begin with. Throughout we will assume that ε and δ are constants and the use of $O(\cdot)$ and $\Omega(\cdot)$ notation in this paper hides all dependence on ε and δ.

Undirected Graphs. Katzir *et al.* [12] implicitly assume the ability to sample from the stationary distribution. They show that in this setting $O\left(\frac{1}{\|\pi\|_2} + d_{\text{avg}}\right)$ samples from the stationary distribution π suffice, where d_{avg} is the average degree. If the graph is connected and a suitable bound on the mixing time exists which is known to the algorithm, $O(t_{\text{mix}} \log n)$ queries suffice to draw one node from a distribution that is close (up to inverse polynomial factors in variation distance) to the stationary distribution using only neighbour queries. This gives an upper bound of $O\left(t_{\text{mix}} \cdot \log n \cdot \left(\frac{1}{\|\pi\|_2} + d_{\text{avg}}\right)\right)$ queries with a neighbour query oracle (Corollary 3.2). In terms of lower bounds, we establish that

- (Theorem 3.5) Any algorithm that has access to random samples from the *stationary distribution* π and outputs a 0.1-accurate estimate of the number of nodes with probability at least 0.99, requires $\Omega\left(\frac{1}{\|\pi\|_2} + d_{\text{avg}}\right)$ samples.
- (Theorem 3.8) Any algorithm that has access to neighbour queries and outputs a 0.1-accurate estimate of the number of nodes with probability at least 0.99, requires $\Omega\left(\frac{1}{\|\pi\|_2} + d_{\text{avg}}\right)$ queries.

We remark that there is a gap between the upper bound and lower bound when considering an oracle with neighbour access. A question left open by our work is whether the multiplicative factor of $t_{\text{mix}} \log n$ is required, or whether a more efficient estimator can be designed.

Directed Graphs. The estimator of Katzir *et al.* is not applicable in the setting when graphs are directed, unless the query model allows in-neighbour queries as well as out-neighbour queries, in which case all results in the undirected setting hold. The reason for this is that even if one did receive a sample drawn from the stationary distribution, it is no longer the case that $\pi_v \propto \deg(v)$, a crucial property exploited by the estimator of Katzir *et al.*. In fact, unless strong assumptions are made on the relationship between the in-degree and the out-degree (*e.g.,* being Eulerian), no simple expression for π in terms of degrees exists.[4]

We provide constructions of graphs that demonstrate that low average degree, rapid mixing and small diameter are not sufficient to design algorithms to estimate the graph size with sublinear (in n) query complexity. In fact, we show that when only given access to a sample from the stationary

[4]We mention that the distribution is closely related to the PageRank distribution. However, the PageRank random walk jumps to a uniformly random node in the graph with a small probability; this is done to avoid problems when the graph is not strongly connected.

distribution, a superpolynomially large sample may be required even for graphs with constant average degree, logarithmic diameter and rapid mixing. The reason for this is that in the case of directed graphs most of the stationary mass can be concentrated on a very small number of nodes.

In order to understand the query complexity when neighbour-queries are allowed, we define a new parameter called *generalized conductance*. Roughly speaking a graph has ε-general conductance ϕ_ε, if every set containing at most $(1-\varepsilon)n$ nodes has at least ϕ_ε fraction of directed edges going out. If a suitable bound on the value of ϕ_ε is known, we show that a simple edge-sampling algorithm outputs an estimate to within relative error ε while using $O(n/\phi_\varepsilon)$ queries (Theorem 4.4) and that this is almost tight, in the sense that any algorithm that outputs an estimate that is even slightly better than ε requires at least $\Omega(n/\phi_\varepsilon)$ queries (Theorem 4.6). The algorithm only requires access to out-neighbours, while the lower bound holds even with respect to an oracle that allows in-neighbour queries, which means that it also applies in the case of undirected graphs.

1.3 Discussion

In terms of improvements to our results, the most interesting question is whether any reasonable subclass of directed graphs are amenable to significantly improved query complexity (ideally sublinear) for the problem of estimating the number of nodes. The constructions in Section 4 show that having low average degree, small diameter and rapid mixing is not enough. For undirected graphs, it is an interesting question whether the extra factor of mixing time t_{mix} must be paid, when only neighbour queries are allowed. It is conceivable that an improved estimator that can handle correlated pairs of nodes can be designed, so as to not waste all but one query for every t_{mix} queries. Finally, it'd be interesting to study the question of estimating other properties of graphs, number of edges, number of triangles, *etc.* in this framework.

The model choices we made reflect the publicly available access to most extant networks; in particular, we were very stringent with accounting—every neighbour query counts as unit cost. Many APIs return the list of neighbours, although in chunks of a fixed size, *e.g.,* 100. It is hard to argue that 100 should be treated as constant in the context of social networks. Nevertheless, if we wanted a list of all followers of Barack Obama, this would still result in a very large number of queries. A natural extension to the query model in this paper is to allow the entire neighbourhood (possibly restricted to the out-neighbourhood in the case of directed graphs) to be revealed at unit cost. Estimators such as the one involving common neighbours of Hardiman and Katzir [11] can be implemented efficiently under such a model. Understanding the query complexity of estimation in these more powerful models is an interesting question.

2 Model and Preliminaries

Graphs. Graphs $G=(V,E)$ considered in this paper may be directed or undirected; typically we assume $|V|=n$ and $|E|=m$, though if there is scope for confusion we use $|V|$ or $|E|$ explicitly. For undirected graphs, for a node $v\in V$, we denote by $N(v)$ its *neighbourhood*, *i.e.,* $N(v)=\{w \mid \{v,w\}\in E\}$, and its degree by $\deg(v) := |N(v)|$. In the case of directed graphs, we denote $N^+(v) := \{w \mid (v,w) \in E\}$ its *out-neighbourhood* and by $\deg^+(v) := |N^+(v)|$ its *out-degree*. Similarly, $N^-(v) := \{u \mid (u,v) \in E\}$ denotes its *in-neighbourhood* and $\deg^-(v) := |N^-(v)|$ its *in-degree*. Furthermore, d_{avg} denotes the average degree, i.e, $\sum_{v\in V}\deg(v)/n$. Whenever there is scope for confusion, we use the notations $\deg_G(u)$, $N_G(v)$, $d_{\text{avg}}(G)$, *etc.* to emphasise that the terms are with respect to graph G.

Random walks in graphs. A discrete-time lazy random walk $(X_t)_{t\geq 0}$ on a graph $G=(V,E)$ is defined by a Markov chain with state space V and transition matrix $P=(p(u,v))_{u,v\in V}$ defined as follows: For every

$u \in V$, $p(u,u) = 1/2$ (*Laziness*). In the undirected setting, for every $v \in N(u)$, $p(u,v) = 1/(2\deg(u))$. In the directed setting, for every $v \in N^+(u)$, $p(u,v) = 1/(2\deg^+(u))$. The transition probabilities can be expressed in matrix form as $P = (I + \mathcal{D}^{-1}A)/2$, where A is the adjacency matrix of G, $\mathcal{D}$ is the diagonal matrix of node degrees (only out-degrees if G is directed), and I is the identity. Let $p^t(u,\cdot)$ denote the distribution over nodes of a random walk at time step t with $X_0 = u$. For the most part, we will consider (strongly) connected graphs. Together with laziness, this ensures that the stationary distribution of the random walk, denoted by π, is unique and given by $\pi P = \pi$. In the undirected case, the form of the stationary distribution is particularly simple, $\pi(u) = \deg(u)/(2|E|)$; furthermore, the random walk is reversible, *i.e.,* $\pi(u)p(u,v) = \pi(v)p(v,u)$. As before, π_G is used to emphasise that the stationary distribution is respect to graph G.

Mixing time. To measure how far $p^t(u,\cdot)$ is from the stationary distribution we consider the *total variation distance*; for distributions μ,ν over sample space Ω the total variation distance is $\|\mu - \nu\|_{\mathsf{TV}} = \frac{1}{2}\sum_{x \in \Omega} |\mu(x) - \nu(x)|$. The *mixing time* of the random walk is defined as $t_{\text{mix}} := \max_{u \in V} \min\{t \geq 1 \mid \left\|p^t(u,\cdot) - \pi\right\|_{\mathsf{TV}} \leq e^{-1}\}$. Although the choice of e^{-1} is arbitrary, it is known that after $t_{\text{mix}} \log(1/\varepsilon)$ steps, the total variation distance is at most ε.

2.1 Query Model

In this section, we formally define the query model that allows us to access the graph. We consider four different neighbour query oracles, $\mathcal{O}$, $\overrightarrow{\mathcal{O}}$, $\overrightarrow{\mathcal{O}}(1)$ and $\overrightarrow{\mathcal{O}}(2)$. We also assume that all oracles have graphs stored as adjacency lists. In the case of directed graphs, there are two adjacency lists for every vertex, one for in-neighbours and one for out-neighbours. No assumption is made regarding the order in which the adjacency lists are stored.

All oracles also make use of labelling functions[5] ; for some space L, a labelling function $\ell : V \to L$ used by an oracle is an injection. We allow ℓ to be defined dynamically. The labelling function we use throughout the paper is the *consecutive labelling function* defined as follows: The label set is $L = \mathbb{N}$. If S denotes all the vertices labelled by the oracle so far, for a new vertex $v \notin S$ picked by the oracle, the label is assigned as follows: if $S = \emptyset$, $\ell(v) = 1$, else $\ell(v) = \max\{\ell(u) \mid u \in S\} + 1$. For these neighbour query oracles, any algorithm can essentially make two types of queries, (i) init, and (ii) (l,i) for undirected graphs or (l,i,etype) for directed graphs, where etype is either in, out. We assume that oracles use a labelling function ℓ.

(i) init: The oracle initialises a set $S := \{v\}$,
where v is chosen to be an arbitrary node in the graph. The different oracles respond as defined below.

- $\mathcal{O}$ responds with $(\ell(v),\deg(v))$ for some $v \in V$. $\overrightarrow{\mathcal{O}}$ responds with $(\ell(v),\deg^+(v))$ for some $v \in V$.
- $\overrightarrow{\mathcal{O}}(1)$ and $\overrightarrow{\mathcal{O}}(2)$ both respond with $(\ell(v),\deg^+(v),\deg^-(v))$.

(ii) (l,i)
or (l,i,etype): $\mathcal{O}$ only responds to query of type (l,i) and the remaining to queries of the type (l,i,etype).

- For query (l,i), if there is $v \in S$ such that $\ell(v) = l$ and $i \leq \deg(v)$, then $\mathcal{O}$ returns $(\ell(u),\deg(u))$, where u is the i^{th} element in the adjacency list of v. The oracle updates $S \leftarrow S \cup \{u\}$. Otherwise it returns null.

[5] We do assume that there is unit cost in sending the label of a node to the algorithm, thus implicitly we may think of every label having a bit-representation that is logarithmic in the size of the graph. However, the bit-length of the label reveals minimal information, which we assume that the algorithm has access to anyway.

- For query (l,i,out), if there is $v \in S$ such that $\ell(v)=l$ and $i \leq \deg^+(v)$, then $\overrightarrow{\mathcal{O}}$ returns $(\ell(u),\deg^+(u))$, while both $\overrightarrow{\mathcal{O}}(1)$ and $\overrightarrow{\mathcal{O}}(2)$ return $(\ell(u),\deg^+(u),\deg^-(u))$ where u is the i^{th} element on the out-neighbour adjacency list of v. The oracles all update $S \leftarrow S \cup \{u\}$. Otherwise, all oracles return null.
- For query (l,i,in), $\overrightarrow{\mathcal{O}}$ and $\overrightarrow{\mathcal{O}}(1)$ always return null. If there exists $v \in S$, such that $\ell(v)=l$ and $i \leq \deg^-(v)$, then $\overrightarrow{\mathcal{O}}(2)$ returns $(\ell(u),\deg^+(u),\deg^-(u))$, where u is the i^{th} element on the adjacency list of in-neighbours of v; otherwise, it returns null. In the case of $\overrightarrow{\mathcal{O}}(2)$, it updates $S \leftarrow S \cup \{u\}$.

In words, $\mathcal{O}$ captures access to undirected graphs such as social networks like Facebook, $\overrightarrow{\mathcal{O}}$ captures directed graphs such as the world wide web, where only out-edges are available, $\overrightarrow{\mathcal{O}}(1)$ captures directed graphs such as fan-networks, where the in-degree but not in-neighbours may be available, and $\overrightarrow{\mathcal{O}}(2)$ captures directed graphs such as Twitter where access is available to both in-edges and out-edges.

It is worth pointing out that a response of null provides no new information to the algorithm. The algorithm only knows that it hadn't received l as a label before or that the degree of the node with label l is strictly smaller than i, things it already knew. This is because the oracle maintains a history of past queries; if this were not the case the algorithm could generate (random) labels and try to find out whether they corresponded to nodes in the graph. However, even for oracles that don't maintain such state explicitly, by choosing ℓ to be collision-resistant hash function, essentially the same behaviour can be achieved.

In some of our proofs, we also allow oracles to return side information. These are denoted by a superscript s, *e.g.,* $\mathcal{O}^s$. Access to oracles with (truthful) side information can only help to reduce query complexity, since an algorithm can choose to ignore any side information it receives. Finally, we say that an algorithm is *sensible* if it does not make a query to which it already knows the answer. It is clear that given any algorithm, there is a sensible algorithm that is at least as good as the original algorithm. The *sensible* algorithm merely simulates the original algorithm and whenever the original algorithm made a query to which the answer was known, the *sensible* algorithm simulates the oracle response.

Finally, we consider the *stationary query oracle*, $\mathcal{O}^\pi$, which when queried returns $(\ell(v),\deg(v))$, where $v \sim \pi$; π is the stationary distribution.

3 Undirected graphs

In this section, we focus on the question of estimating the number of nodes in undirected graphs. We show that the results obtained by Katzir *et al.* [12] are essentially tight, up to a factor of the mixing time and polylogarithmic terms in n. This suggests that rapid mixing is a critical condition for being able to estimate the size of the graph. We begin by discussing very simple examples which show that this is indeed the case.

Some Simple Observations First we construct two simple graphs, one having twice as many nodes as the other. No algorithm can distinguish between these two graphs with probability greater than 3/4 (say) unless $\Omega(m)$ queries to the oracle $\mathcal{O}$ are made. The graphs are shown in Figure 1. The first graph is simply $G_{n,p}$ with each node being connected to an additional node with degree 1. The second graph is identical, except that for one randomly chosen node, instead of this additional edge to a degree 1 node, it is connected to a copy of $G_{n,p}$. Suppose $p > \log(n)/n$ so that the graphs are connected with high probability and that $m = \Theta(n^2 p)$.

We don't provide a formal proof here, which can be given along the lines of the proof of Theorem 3.5. We outline a sketch here. Assuming the oracle $\mathcal{O}$ returns some $v \in V$ in the part that is $G_{n,p}$ with an additional edge coming out of each node, how long will it take before we discover the edge to the copy

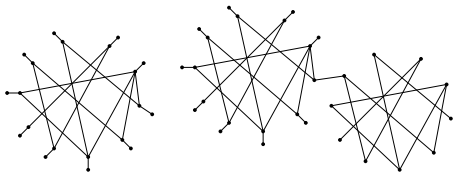

Figure 1: Example showing that $\Omega(m)$ samples might be required if no assumption like rapid mixing is made.

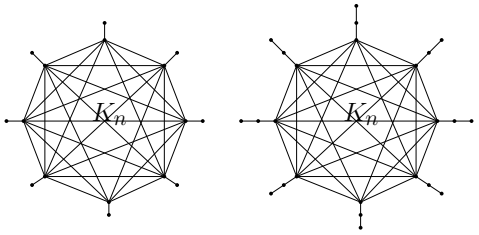

Figure 2: The *sun* graph and the *bright sun*.

of $G_{n,p}$, if it exists? Since we made no assumption on the order in which neighbours are stored in an adjacency list, we have to query at least a constant fraction of neighbours of a given node to discover, with constant probability, the edge that leads out of the $G_{n,p}$. However, most of these will just lead us to a single vertex. Thus, unless we explore at least a constant fraction of the edges of a constant fraction of the vertices, we don't stand a significant chance of finding the edge that leads to the copy of the $G_{n,p}$. Thus, $\Omega(m)$ queries to the oracle are required.

Clearly, if we're in the second case, that is the graph with an additional copy of $G_{n,p}$ connected to the earlier one by a single edge, a random walk on this graph does not mix rapidly. What this example shows is that if one wishes to do asymptotically better than breadth-first search, one needs additional assumptions on the graph, such as rapid mixing.

3.1 Results of Katzir *et al.*

In this section, we discuss the result of Katzir *et al.* [12] regarding estimating the number of nodes in a graph. Though, they don't discuss this formally, their method essentially boils down to having access to the stationary query oracle $\mathcal{O}^\pi$.

For the sake of completeness, we outline the estimator of Katzir *et al.* here. Let

$$X = \{(\ell(x_1),\deg(x_1)),(\ell(x_2),\deg(x_2)),...,(\ell(x_r),\deg(x_r))\}$$

be drawn from the stationary query oracle $\mathcal{O}^\pi$. Let $\Psi_1 = \sum_{i=1}^r \deg(x_i)$ and $\Psi_{-1} = \sum_{i=1}^r 1/\deg(x_i)$; let $C = \sum_{i \neq j} \mathbb{1}(\ell(x_i) = \ell(x_j))$ denote the random variable counting the number of collisions. Then, it is fairly straightforward to see that $\mathbb{E}[\Psi_1 \Psi_{-1}] = r + 2n\binom{r}{2}\|\pi\|_2^2$ and $\mathbb{E}[C] = 2\binom{r}{2}\|\pi\|_2^2$. Katzir *et al.* use the the following as an estimator for the number of nodes: $\widehat{n} := \frac{\Psi_1 \Psi_{-1} - r}{C}$ They prove the following result.

Theorem 3.1 (Katzir et al. [12])**.** *Let $\varepsilon > 0$ and $\delta > 0$. Suppose that the number of samples is $r \geq 1 + \frac{32}{\varepsilon^2 \delta} \max\left\{\frac{1}{\|\pi\|_2}, d_{\text{avg}}\right\}$, where $d_{\text{avg}} = 2m/n$. Then, $\mathbb{P}[|\widehat{n} - n| \geq \varepsilon n] \leq \delta$.*

Given a bound T on t_{mix} and access to oracle $\mathcal{O}$, the stationary query oracle $\mathcal{O}^\pi$ can be approximately simulated. (We assume that the graph is connected.) We simply perform a random walk on the graph for $T\log(1/\rho)$ steps, if a sample from a distribution at most ρ far from π in total variation distance is desired. Using the above theorem, we get the following straightforward corollary. The proof of the corollary follows by simulating s random walks for $O(T\log s)$ time steps (queries) ensuring that each random walk has a total variation distance of s^{-2} from π. Using the above theorem, we get the following straightforward corollary.

Corollary 3.2. *Let $\varepsilon > 0$ and $\delta > 0$. For a connected, undirected graph, $G = (V,E)$ let T be such that $t_{\text{mix}} \leq T$. Then there exists an algorithm that given T and access to oracle $\mathcal{O}$, outputs $\widehat{n}$, such that,*

$\mathbb{P}[|\widehat{n}-n|\geq \varepsilon n]\leq \delta$. *Furthermore, the number of queries made by the algorithm to* $\mathcal{O}$ *is* $O(T s \log s)$, *where* $s = O\left(\frac{1}{\varepsilon^2\delta}\cdot \max\left\{\frac{1}{\|\pi\|_2}, d_{\text{avg}}\right\}\right)$.

The key question is, are these bounds tight? In Section 3.2, we give much stronger lower bounds, where we show that for any valid degree sequence $\mathbf{d}=(d_1,\ldots,d_n)$ on n nodes, there exists a sequence $\mathbf{d}'$ on $2n$ nodes, such that any algorithm with access to a stationary oracle for either sequence $\mathbf{d}$ or $\mathbf{d}'$, cannot distinguish between the two unless it makes $\Omega\left(\frac{1}{\|\pi\|_2}+d_{\text{avg}}\right)$ queries. Here, we discuss simple instances where the bound in Theorem 3.1 is tight; we don't give formal proofs which are relatively straightforward.

- Let G_1 and G_2 be d-regular graphs on n and $2n$ nodes respectively, for $d<n$. Thus, samples from the stationary distribution are essentially just uniformly chosen random nodes in the graph. As the degrees are identical, they reveal no additional information. Thus, the algorithm has to query until a collision is observed, which requires $\Omega(\sqrt{n})$ queries. Clearly for these graphs $\frac{1}{\|\pi\|_2}=\Theta(\sqrt{n})$.
- We define the *sun* graph as a K_n with an additional edge out of each vertex to a node with degree 1. The *bright sun* graph is the same as the sun graph, except that there is a path of length 2 coming out of each vertex, rather than just an edge (see Figure 2). Under the stationary oracle model, $\Omega(n)$ queries are required before any degree 1 or 2 nodes are returned since the total probability on these nodes is $O(1/n)$. As can be seen for these graphs $d_{\text{avg}}=\Theta(n)$.

The above examples show that there are graphs for which the bound of Katzir *et al.* is tight. In this paper, we show a significantly stronger statement—given any degree sequence $\mathbf{d}$, there are graphs for which the bound is almost tight.

3.2 Lower Bounds in the Stationary Query Model

In this section, we show that for any undirected, connected graph $G=(V,E)$, there exists an undirected, connected graph $\widetilde{G}$, such that any algorithm which has access to either the oracle $\mathcal{O}^\pi(G)$ or $\mathcal{O}^\pi(\widetilde{G})$, cannot distinguish between the two with significant probability without making a large number of queries. Thus, it cannot output an estimate $\widehat{n}$ satisfying $|\widehat{n}-|V||<|V|/2$ with probability $2/3$.

Lemma 3.3. *Let* $G=(V,E)$ *be an undirected, connected graph with* $|V|=n$ *and* $|E|\geq n$. *Then there exists a graph* $\widetilde{G}=(\widetilde{V},\widetilde{E})$, *with* $|\widetilde{V}|=2n$, *such that any algorithm given access to either* $\mathcal{O}^\pi(G)$ *or* $\mathcal{O}^\pi(\widetilde{G})$ *with equal probability, cannot distinguish between the two with probability greater than* $\frac{2}{3}$, *unless it makes at least* $\Omega\left(\frac{1}{\|\pi_G\|_2}\right)$ *queries. As a consequence, no algorithm can output* $\widehat{n}$ *satisfying* $|\widehat{n}-n^*|<n^*/2$ *w.p. at least* $2/3$, *where* $n^*=n$ *if the chosen graph is* G *and* $2n$ *if it is* $\widetilde{G}$.

Proof. We construct $\widetilde{G}=(\widetilde{V},\widetilde{E})$ by taking $G=(V,E)$ and an identical copy, denoted $G'=(V',E')$, and connecting them arbitrarily in such a way that the degrees of the nodes remain unchanged. We note that this can always be done. Since $|E|\geq n$, the graph G is not a tree, and hence there is an edge $\{i,j\}\in E$ such that removing it does not disconnect G. Let $\{i',j'\}$ be the corresponding edge in the copy of G, then removing $\{i,j\}$ and $\{i',j'\}$ and adding the edges $\{i,j'\}$ and $\{i',j\}$ gives $\widetilde{G}$.

For some $s\in\mathbb{N}$, let $(X_t)_{t=1}^s$ be s independent samples drawn from π_G. Similarly, let $(\widetilde{X}_t)_{t=1}^s$ be independent samples from $\pi_{\widetilde{G}}$. We define a coupling between $(X_t)_{t=1}^s$ and $(\widetilde{X}_t)_{t=1}^s$. Suppose $X_t=i$, then we define $\widetilde{X}_t=i$ with probability $\frac{1}{2}$ and $\widetilde{X}_t=i'$ (the copy of i in G') with probability $\frac{1}{2}$. This ensures that

$\widetilde{X}_t$ is drawn from the distribution $\pi_{\widetilde{G}}$. Note that the output of the oracle $\mathcal{O}^\pi(G)$ (respectively $\mathcal{O}^\pi(\widetilde{G})$) is $(\ell(X_t),\deg_G(X_t))$ (respectively $(\ell(\widetilde{X}_t),\deg_{\widetilde{G}}(\widetilde{X}_t))$).

We assume the oracle simply uses a consecutive labelling scheme, *i.e.,* when X_t is drawn from π, if there is no $t'<t$, such that $X_{t'}=X_t$, then $\ell(X_t)=\max\{\ell(X_{t'}) \mid t'<t\}+1$. Note that this ensures that unless there exist $t\neq t'$ and $t,t'\leq s$, such that $X_t=X_{t'}$, the outputs of oracles $\mathcal{O}^\pi(G)$ and $\mathcal{O}^\pi(\widetilde{G})$ would be identical.

Finally, define

$$C_s=\sum_{t'<t\leq s}\mathbb{1}(X_t\neq\widetilde{X}_{t'})$$

The expectation of C_s is $\mathbb{E}[C_s]=\binom{s}{2}\|\pi_G\|_2^2$. Then by Markov's inequality, for a suitably chosen constant c, if $s\leq\frac{c}{\|\pi_G\|_2}$, $\mathbb{P}[C\geq 1]\leq 1/6$. Thus, with probability at least $\frac{5}{6}$, the outputs of oracles $\mathcal{O}^\pi(G)$ and $\mathcal{O}^\pi(\widetilde{G})$ can be coupled perfectly. □

Lemma 3.4. *Let $G=(V,E)$ be an undirected, connected graph with $|V|=n$ and $|E|\geq n$. Then there exists a graph $\widetilde{G}=(\widetilde{V},\widetilde{E})$, with $|\widetilde{V}|=2n$, such that any algorithm given access to either $\mathcal{O}^\pi(G)$ or $\mathcal{O}^\pi(\widetilde{G})$ with equal probability, cannot distinguish between the two with probability greater than $\frac{2}{3}$, unless it makes at least $\Omega(d_{\text{avg}}(G))$ queries. As a consequence, no algorithm can output $\widehat{n}$ satisfying $|\widehat{n}-n^*|<n^*/2$ w.p. at least 2/3, where $n^*=n$ if the oracle chosen corresponds to G and $n^*=2n$ if it corresponds to $\widetilde{G}$.*

Proof. We construct $\widetilde{G}=(\widetilde{V},\widetilde{E})$ by taking $G=(V,E)$ and a 3-regular expander on n nodes, denoted $G'=(V',E')$, and connecting them arbitrarily in such a way that the degrees of the nodes remain unchanged. (If n is odd, we can use set $|V'|=|V|+1$.) We note that this can always be done. Since $|E|\geq n$, the graph G is not a tree, and hence there is an edge $\{i,j\}\in E$ such that removing it does not disconnect G. Let $\{i',j'\}$ be the corresponding edge in the copy of G, then removing $\{i,j\}$ and $\{i',j'\}$ and adding the edges $\{i,j'\}$ and $\{i',j\}$ gives $\widetilde{G}$.

For $s\in\mathbb{N}$, let $(X_t)_{t=1}^s$ be s independent samples drawn from π_G; similarly $(\widetilde{X}_t)_{t=1}^s$ a sample drawn from $\pi_{\widetilde{G}}$. We define a coupling between $(X_t)_{t=1}^s$ and $(\widetilde{X}_t)_{t=1}^s$ as follows. Let $p=\pi_{\widetilde{G}}(V')$, the probability that a node sampled according to $\pi_{\widetilde{G}}$ is one of the newly created ones. Let $Z_t=1$ with probability p and 0 otherwise. Then we have $\widetilde{X}_t=X_t$ if $Z_t=0$, otherwise $\widetilde{X}_t$ is chosen uniformly at random from V', the newly added nodes. To see that this is a valid coupling, observe that for any $u\sim\pi_{\widetilde{G}}$, conditioned on $u\notin V'$, u is distributed according to π_G.

Note that the oracle $\mathcal{O}^\pi(G)$ (respectively $\mathcal{O}^\pi(\widetilde{G})$) returns $(\ell(X_t),\deg_G(X_t))$ (respectively $(\ell(\widetilde{X}_t),\deg_{\widetilde{G}}(X_t))$). We assume that the oracles use a consecutive labelling function, *i.e.,* every new node is given the smallest as yet unused natural number.

Let $C=\sum_{t=1}^s Z_t$, then $\mathbb{E}[C]=ps$. Note that $p=3|V'|/(2|E|+3|V|')\leq 3|V|/(2|E|)=3/d_{\text{avg}}(G)$. Thus, for a suitably chosen constant c, if $s\leq c\cdot d_{\text{avg}}(G)$, by Markov's inequality $\mathbb{P}[C\geq 1]\leq 1/6$. Note that conditioned on $C=0$, the outputs of the oracles $\mathcal{O}^\pi(G)$ and $\mathcal{O}^\pi(\widetilde{G})$ can be coupled perfectly. Hence, no algorithm can distinguish between G and $\widetilde{G}$ using fewer than $c\cdot d_{\text{avg}}(G)$ queries with probability greater than or equal to 2/3. □

As a consequence of Lemma 3.3 and Lemma 3.4, we get the following theorem.

Theorem 3.5. *Given an undirected, connected graph $G=(V,E)$, there exist graphs $\widetilde{G}$, $\bar{G}$ with $2|V|$ nodes and a constant $p<1$, such that any algorithm that is given access to one of three oracles $\mathcal{O}^\pi(G)$, $\mathcal{O}^\pi(\widetilde{G})$ and $\mathcal{O}^\pi(\bar{G})$, chosen with equal probability, requires $\Omega\left(\frac{1}{\|\pi_G\|_2}+d_{\text{avg}}(G)\right)$ queries to distinguish between them with probability at least p. As a consequence, any algorithm that outputs $\widehat{n}$, such that $|\widehat{n}-n^*|<n^*/2$ requires at least $\Omega\left(\frac{1}{\|\pi_G\|_2}+d_{\text{avg}}(G)\right)$ queries, where $n^*=n$ if the graph is G and $n^*=2n$ if the graph is either $\widetilde{G}$ or $\bar{G}$.*

3.3 Oracle Sampling from the Neighbour Query Model

In this section, we show that with access to the oracle $\mathcal{O}(G)$, any algorithm that predicts the number of nodes in a graph G to within a small constant fraction requires $\Omega\left(\frac{1}{\|\pi_G\|_2}+d_{\text{avg}}\right)$ queries. For proving the lower bounds we use graphs generated according to the *configuration model* [2]. A vector $\mathbf{d}=(d_1,d_2,...,d_n)$ is said to be *graphical* if there exists an undirected graph on n nodes such that vertex $i\in[n]$ has degree d_i. We briefly describe here how graphs are generated in the configuration model:

1. Create disjoint sets W_i, for $i\in\{1,...n\}$, with $|W_i|=d_i$. The elements of W_i are called *stubs*.
2. Create a uniform random maximum matching in the set $\bigcup_{i=1}^{n} W_i$ (note that $\sum_{i=1}^{n} d_i$ must be even since $\mathbf{d}$ is graphical).
3. For a stub edge $\{x,y\}$ in the matching, such that $x\in W_i$ and $y\in W_j$, the edge $\{i,j\}$ is added to the graph.

The above procedure creates a graph where vertex i has degree exactly d_i. However, the graph may not be simple, *i.e.,* it may have multiple edges and self-loops. Also, this procedure does not necessarily produce a uniform distribution over graphs having degree sequence $\mathbf{d}$. The expected number of multi-edges and self-loops in the graph is for many interesting graphs only a small fraction and in any graph with bounded degree their expected number is a constant. We use $G\sim\mathcal{G}(\mathbf{d})$ to denote that a graph G was generated in the configuration model with degree sequence $\mathbf{d}$.

Recall the definition of a *sensible* algorithm as one that never makes a query to which it already knows the answer. A *sensible* algorithm has the following behaviour: (i) for every $u\in V$ and $i\leq\deg(u)$ ($\deg^+(u)$ and $\deg^-(u)$, respectively) it makes the query $(\ell(u),i)$ ($(\ell(u),i,\mathsf{in})$ and $(\ell(u),i,\mathsf{out})$, respectively) at most once, and (ii) it never queries a $(\ell(v),i)$ if it has not received $\ell(v)$ as a valid label or if $i\notin[1,\deg(v)]$. Note that there exists for any algorithm a sensible implementation which needs at most as many queries as the original algorithm. For technical reasons, in the proof of Lemma 3.6, we use an oracle $\mathcal{O}^s$ with *s*ide-information as an extension of $\mathcal{O}$: $\mathcal{O}^s$ returns, upon query, exactly the same information as $\mathcal{O}$, but can add additional truthful information. In particular, we allow the oracle when queried $(\ell(v),i)$ to not only return the corresponding node $(\ell(u),\deg(u))$, where u is the i^{th} element in the adjacency list of v, but also the index, say j, in the adjacency list of u which corresponds to v. Clearly, any sensible algorithm wouldn't query $(\ell(u),j)$ after querying $(\ell(v),i)$.

Lemma 3.6. *Let $\varepsilon>0$. Let $\boldsymbol{d}=(d_1,...,d_n)$ be an arbitrary graphical sequence and let $D:=\sum_{i=1}^{n} d_i$. Let $G\sim\mathcal{G}(d)$, and let $\widetilde{G}$ be an arbitrary graph with degree sequence $\boldsymbol{d}$. There exists an implementation of an oracle $\mathcal{O}^s(G)$ (with side-information) such that if $(\ell(X_t),\deg_G(X_t))_{t=1}^{T}$ is the sequence of responses to a* sensible *algorithm, where*

$$T:=\min\Big\{\min\{\tau\geq 0 \mid \textit{all neighbours of all known nodes are disclosed using } \tau \textit{ queries}\},\frac{\varepsilon}{16}\cdot\sqrt{D}\Big\},$$

and if $(\ell(\widetilde{X}_t),\deg_{\widetilde{G}}(\widetilde{X}_t))_{t=1}^{T}$ is the sequence returned by oracle $\mathcal{O}^\pi(\widetilde{G})$, then there exists a coupling so that the sequences $((\ell(X_t),\deg_G(X_t)))_{t=1}^{T}$ and $((\ell(\widetilde{X}_t),\deg_{\widetilde{G}}(\widetilde{X}_t)))_{t=1}^{T}$ are identical with probability at least $1-\varepsilon$.

Remark 3.7. *Graphs $G\sim\mathcal{G}$ are not necessarily connected. Thus, Lemma 3.6 uses stopping time T where all edges in the connected component involving the starting node, i.e., the node sent in response*

to the init *are uncovered by the algorithm, either through query responses or through side information. Clearly, at this point there is nothing left of the algorithm to do except return the number of observed nodes as its estimate, which can be arbitrarily far off.*

Proof of Lemma 3.6 . We consider the following implementation of $\mathcal{O}^s(G)$, which generates the graph $G\sim\mathcal{G}(\mathbf{d})$ on the fly as it receives the queries. We assume that $\mathcal{O}^s(G)$ uses a consecutive labelling function, *i.e.,* whenever it chooses a new node, it picks the smallest as yet unused natural number as the label. In response to the init, it picks $X_1\sim\pi$, sets $\ell(X_1)=1$ and returns $(\ell(X_1),\deg_G(X_1))$.

It starts by creating sets $W_u=\{u_1,u_2,...,u_{d_u}\}$ for $u=1,...,n$. Let $\mathcal{F}_t$ be the filtration corresponding to all the random choices made including the answer to the t^{th} query (counting the init as the first one). Let (l,i) be the $(t+1)^{th}$ query made by the algorithm, let $v\in V$, such that $\ell(v)=l$. Since the algorithm is *sensible*, it must be that (l,i) is queried for the first time and that the algorithm never learned, through side-information of the oracle, to which node the stub v_i connects; this means $\mathcal{O}^s(G)$ has not fixed the choice of partner for v_i yet.

The oracle chooses a random unmatched stub, say $v'_{i'}$ and returns $(\ell(v'),\deg_G(v'))$ as well as the *side information* i', indicating that v is the $(i')^{th}$ neighbour of v'.

Let p^t denote the distribution over choice of v' made by $\mathcal{O}^s(G)$ before responding to the $(t+1)^{th}$ query. Let $\delta_u(t)$ denote number of edges of u that have been already disclosed up to the answering of the t^{th} query. This also includes possible side-information. Note that $\sum_u\delta(u)=2(t-1)$, since every query other than the init reveals information about the two end-points of the disclosed edge. Thus, $p^t(u)=(d_u-\delta_u(t))/(D-2(t-1))$. Consider the variation distance between p^t and π,

$$\begin{aligned}\left\|\pi-p^t\right\|_{\mathsf{TV}}&=\frac{1}{2}\sum_{u\in V}\left|\frac{d_u}{D}-\frac{d_u-\delta_u(t)}{D-2(t-1)}\right|=\frac{1}{2}\sum_{u\in V}\left|\frac{-2(t-1)d_u+D\delta_u(t)}{D(D-2t+2)}\right|\\&\le\frac{1}{2}\sum_{u\in v}\frac{2(t-1)d_u}{D(D-2t+2)}+\frac{1}{2}\sum_{u\in V}\frac{D\delta_u(t)}{D(D-2t+2)}\\&\le\frac{1}{2}\cdot\frac{2(t-1)}{D-2t+2}+\frac{1}{2}\cdot\frac{2(t-1)}{D-2t+2}=\frac{2t-2}{D-2t+2}\end{aligned}$$

In the final step above, we used the fact that $\sum_u d_u=D$ and $\sum_u\delta_u(t)=2t-2$. Note that if the stopping condition in the definition of T has not occurred the above calculations are valid. Also, we know that $T\le\frac{\varepsilon}{16}\sqrt{D}$, thus summing up the variation distance $\left\|\pi-p^t\right\|_{\mathsf{TV}}$ for $t\in\{1,...,T\}$ and observing that $\left\|p^0-\pi\right\|_{\mathsf{TV}}=0$, since the response to the init is chosen from the stationary distribution, we get the desired result. □

Theorem 3.8. *Let $\boldsymbol{d}=(d_1,...,d_n)$ be a graphical vector satisfying $\min_i d_i\ge 3$. Then there exists a graphical $\widetilde{\boldsymbol{d}}=(\widetilde{d}_1,...,\widetilde{d}_n,\widetilde{d}_{n+1},...,\widetilde{d}_{2n})$ with $\min_i\widetilde{d}_i\ge 3$, such that for $G\sim\mathcal{G}(\boldsymbol{d})$ and $\widetilde{G}\sim\mathcal{G}(\widetilde{\boldsymbol{d}})$, there exists $c>0$, such that any algorithm with access to one of two oracles $\mathcal{O}(G)$ or $\mathcal{O}(\widetilde{G})$ chosen equal probability, cannot distinguish between the two with probability greater than $1-c$ unless it makes $\Omega\left(\frac{1}{\|\pi_G\|_2}+d_{\mathrm{avg}}(G)\right)$ queries to the oracle.*

Proof. First we notice that $\frac{1}{\|\pi_G\|_2}=O(\sqrt{D})$ and $d_{\mathrm{avg}}(G)=O(\sqrt{D})$ if $D:=\sum_{i=1}^n d_i$.

We first prove the lower bound of $\Omega(1/\|\pi_G\|_2)$. We set $\widetilde{\mathbf{d}}=(d_1,...,d_n,d_1,...,d_n)$ as in the proof of Lemma 3.3. Let G_1 and G_2 be an arbitrary graphs on n and $2n$ vertices with degree sequences $\mathbf{d}$ and $\widetilde{\mathbf{d}}$ respectively. Let $\mathcal{E}_1$ (resp. $\mathcal{E}_2$) be the event that the coupling of the sequences output by $\mathcal{O}^s(G)$ (resp. $\mathcal{O}^s(\widetilde{G})$) and $\mathcal{O}^\pi(G_1)$ (reps. $\mathcal{O}^\pi(G_2)$) holds for the first T queries using Lemma 3.6. Let $\mathcal{E}'_1$ and $\mathcal{E}'_2$ be the events that the stopping condition in the coupling in Lemma 3.6 happens due to $\varepsilon\sqrt{D}/16$ queries being made. Note that because of the condition $\min_i d_i\ge 3$ (and $\min_i\widetilde{d}_i\ge 3$) the graph $G\sim\mathcal{G}(\mathbf{d})$ (and

$\widetilde{G} \sim \mathcal{G}(\widetilde{\mathbf{d}})$) is connected with high probability. Thus, as $n \to \infty$, $\mathbb{P}[\mathcal{E}_1'] \to 1$ (the same for $\mathbb{P}[\mathcal{E}_2']$). Finally, let $\mathcal{E}_3$ be the event that the coupling in Lemma 3.3 holds. By choosing ε in Lemma 3.6 appropriately, we can bound the probability that at least one of the events $\mathcal{E}_1$, $\mathcal{E}_2$, $\mathcal{E}_1'$, $\mathcal{E}_2'$, and $\mathcal{E}_3$ does not occur, by some constant $c < 1$. Thus, we get the require result.

The lower bound of $\Omega(d_{\text{avg}}(G))$ follows similarly using Lemma 3.4 instead of Lemma 3.3. □

4 Directed Graphs

In this section, we consider the query complexity of estimating the number of nodes in directed graphs. We first observe that estimating n using the approach of Katzir *et al.* [12] is not possible since the stationary distribution of a node is in general not proportional to its degree. Another obstacle is that the stationary distribution of a node can be exponentially small as the graphs in Figures 3 and 4 illustrate. In particular, it takes an exponentially large sample drawn from the stationary distribution to distinguish between the line graph of Figure 3 on n nodes and the line graph on $2n$ nodes, since the probability mass of the additional nodes is $2^{-\Omega(n)}$. It is also not very difficult to show that even with access to one of the two oracles $\overrightarrow{\mathcal{O}}$ or $\overrightarrow{\mathcal{O}}(1)$, $\Omega(n)$ queries are required to distinguish the line graph on n vertices from the line graph on $2n$ vertices.

As the example of the line graph reveals, unlike in the undirected case, rapid mixing and low average degree are not sufficient conditions to design a good estimator of the number of nodes using sublinear number of queries. The line graph shows that in the directed case, rapid mixing does not imply short (directed) diameter. One might hope that if one throws small diameter into the mix, in addition to low average degree and rapid mixing, a better estimator could be designed. In Section 4, we show that this is not the case. The problem of estimation remains stubbornly hard, and $\Omega(n)$ queries to the oracle $\overrightarrow{\mathcal{O}}$, and $2^{\Omega(n)}$ queries to the stationary query oracle are required to achieve a good estimate of the number of nodes.

These observations suggest that exploring the graph, *e.g.,* through breadth-first search, is much faster than sampling from the stationary distribution. The question of interest is whether there is a property, satisfied by graphs of interest, which yields a query complexity better than $\Omega(m)$. We answer this positively in Section 4.1, where we introduce a parameter that generalises the conductance ϕ_ε and give almost tight bounds on the number of queries required to estimate n up to an ε relative error. Our Algorithm EdgeSampling takes this parameter as an input and terminates after $O(n/\phi_\varepsilon)$ queries which can be much smaller than the sample complexity of breadth-first search.

Before delving into these results, it is worth pointing out that results in Section 3 can be used if access to $\overrightarrow{\mathcal{O}}(2)$ is provided. In this case, we can simply treat the graph as being undirected. However, it is still interesting to understand whether the distinction between in-neighbours and out-neighbours allows one to design better estimators. At present, we are unaware of any graphs where this might be the case.

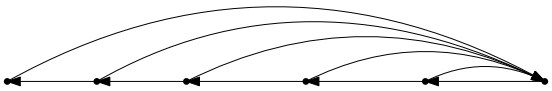

Figure 3: The Line graph on 6 nodes.

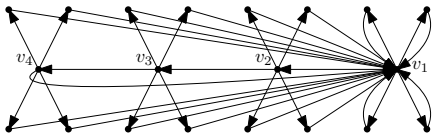

Figure 4: The graph Comet(20,4).

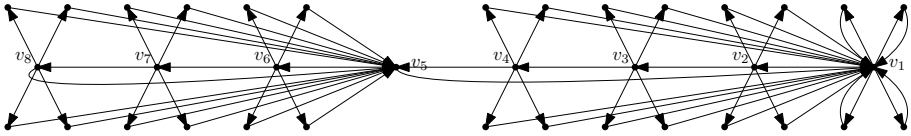

Figure 5: The graph DoubleComet(40,8). DoubleComet($2n$,$2k$) consists of two copies of Comet(n,k) connected in the following way. First we remove all directed neighbours of v_{k+1} as well as the edge (v_k,v_1). Then we add (v_k,v_{k+1}), (v_{k+1},v_{k+2}) as well as (v_{k+1},v_1).

The Comet Graph The *Comet graph*, Comet(n,k) is constructed as follows. Assume that k divides n. There is a directed cycle on the vertices $v_1,v_2,...,v_k$, with edges (v_i,v_{i+1}) for $1\leq i<k$ and (v_k,v_1). We denote these k vertices as *centres*. For every $\ell\in[k]$, there is a directed star $S_\ell=\{(v_\ell,v_{\ell,j}):j\in[n/k-1]\}$ with centre in v_ℓ of degree $n/k-1$. For each leaf $v_{\ell,j}$ in star S_ℓ with $\ell\in[k],j\in[n/k-1]$, there is a directed edge to the first star centre v_1, that is, $\{(v_{\ell,j},v_1):\ell\in[1,k],j\in[n/k-1]\}$. We write v_ℓ^G and $v_{\ell,j}^G$ to emphasise that the nodes belong to graph G.

In the following lemma we obtain bounds on stationary distribution of the nodes, which will allow us to obtain a bound on the query complexity w.r.t. $\mathcal{O}^\pi$ (see Theorem 4.2).

Lemma 4.1. *Let n be a multiple of k. Let π be the stationary distribution of the nodes in Comet(n,k). Then, for $G=$ Comet(n,k) and $\widetilde{G}=$ Comet($2n$,$2k$) it holds that $\left\|\pi_G-\pi_{\widetilde{G}}\right\|_{\mathsf{TV}}=O((\frac{k}{n})^{k-1})$.*

Proof. Note that for $u\in V$ we have $\pi(v)=\sum_{u\in V}p(u,v)\pi(u)$. Let $\ell\in[k]$, $u\in N^+(v_\ell)\setminus\{v_1\}$ and $d=n/k$. We have that

$$\pi(u)=\pi(v_\ell)/2d+\pi(u)/2. \tag{1}$$

This implies that $\pi(v_\ell)=d\pi(u)$ for every $u\in N^+(v_\ell)$. In particular, for $u=v_{\ell+1}$ (consider $v_{k+1}=v_1$) it follows that

$$\pi(v_{\ell+1})=\pi(v_\ell)/d. \tag{2}$$

Therefore, $\pi(v_\ell)=\pi(v_1)/d^{\ell-1}$. Furthermore,

$$\pi(v_{\ell,j})=\pi(v_{\ell+1})=\pi(v_1)/d^\ell \tag{3}$$

for every $\ell\in[k]$, $j\in[d-1]$. Summing up everything,

$$\begin{aligned}
1 &= \sum_{\ell=1}^{k}\left(\pi(v_\ell)+\sum_{j=1}^{d-1}\pi(v_{\ell,j})\right)\\
&= \sum_{\ell=1}^{k}\frac{\pi(v_1)}{d^{\ell-1}}+\sum_{\ell=1}^{k}\sum_{j=1}^{d-1}\frac{\pi(v_1)}{d^{\ell}}\\
&= \pi(v_1)\left(\frac{1-1/d^k}{1-1/d}+(d-1)\frac{1/d-1/d^{k+1}}{1-1/d}\right),\\
&= \pi(v_1)\left(\frac{2d-1-\frac{2d+1}{d^k}}{d-1}\right).
\end{aligned}$$

Rearranging the terms yields

$$\pi(v_1)=\frac{d-1}{2d-1-\frac{2d+1}{d^k}}. \tag{4}$$

We proceed by establishing a bound on $\left\|\pi_G-\pi_{\widetilde{G}}\right\|_{\mathsf{TV}}$.

$$\begin{aligned}
\left\|\pi_G-\pi_{\widetilde{G}}\right\|_{\mathsf{TV}} &= \frac{1}{2}\sum_{\ell=1}^{k}\left|\pi(v_\ell^G)-\pi(v_\ell^{\widetilde{G}})\right|+\frac{1}{2}\sum_{\ell=1}^{k}\sum_{j=1}^{d-1}\left|\pi(v_{\ell,j}^G)-\pi(v_{\ell,j}^{\widetilde{G}})\right|\\
&\quad+\frac{1}{2}\sum_{\ell=k+1}^{2k}\left|0-\pi(v_\ell^{\widetilde{G}})\right|+\frac{1}{2}\sum_{\ell=k+1}^{2k}\sum_{j=1}^{d-1}\left|0-\pi(v_{\ell,j}^{\widetilde{G}})\right|\\
&\le \sum_{\ell=1}^{k}\left|\pi(v_\ell^G)-\pi(v_\ell^{\widetilde{G}})\right|+\sum_{\ell=k+1}^{2k}\left|\pi(v_\ell^{\widetilde{G}})\right|\\
&\le \left|\pi(v_1^G)-\pi(v_1^{\widetilde{G}})\right|\sum_{\ell=1}^{k}\frac{1}{d^{\ell-1}}+\sum_{\ell=k+1}^{2k}\frac{1}{d^{\ell-1}}\\
&\le 2\left|\pi(v_1^G)-\pi(v_1^{\widetilde{G}})\right|+\frac{\frac{1}{d^k}}{1-\frac{1}{d}}.
\end{aligned}$$

By triangle inequality,

$$\left|\pi(v_1^G)-\pi(v_1^{\widetilde{G}})\right|\le\left|\pi(v_1^G)-\frac{d-1}{2d-1}\right|+\left|\pi(v_1^{\widetilde{G}})-\frac{d-1}{2d-1}\right|\le 2\left|\frac{d-1}{2d-1-\frac{2d+1}{d^k}}-\frac{d-1}{2d-1}\right|\le 4\frac{2d+1}{d^k}.$$

Putting everything together yields $\left\|\pi_G-\pi_{\widetilde{G}}\right\|_{\mathsf{TV}}=O\left(\left(\frac{k}{n}\right)^{k-1}\right)$.

□

To establish the mixing time of Comet(n,k), we will make use of the result of Levin, Peres and Wilmer [14] which relates the mixing time of a Markov chain to the probability that two copies of the chain meet.

Theorem 4.2. *Let n be a multiple of k. Then Comet(n,k) has mixing time $t_{\text{mix}}=O(1)$ and diameter k. Furthermore, any algorithm requires at least $\Omega((\frac{n}{k})^{k-1})$ queries to $\mathcal{O}^\pi$ and $\Omega(n)$ queries to $\overrightarrow{\mathcal{O}}$ to distinguish between $G=$ Comet(n,k) and $\widetilde{G}=$ Comet$(2n,2k)$.*

Proof. Observe that the diameter of the graph is indeed k. Let $(X_t)_{t\geq 0}$, $(Y_t)_{t\geq 0}$ be random walks on Comet(n,k) with $X_0=x$ and $Y_0=y$, where $x,y\in V$. Whenever $X_s=Y_s$, then we can couple $X_t=Y_t$ for $t\geq s$. Let $T=\min\{t\geq 0 \mid X_t=Y_t\}$. Starting from any vertex $u\in V$, with constant probability at least $p>0$, the random walk is at v_1 after 2 steps, *i.e.,* $\mathbb{P}[X_2=v_1 \mid X_0=u]\geq p$ for $u\in V$. In particular, using independence, $\mathbb{P}[X_2=Y_2 \mid X_0=u,Y_0=v]\geq \mathbb{P}[X_2=v_1 \mid X_0=u]\cdot\mathbb{P}[Y_2=v_1 \mid Y_0=v]\geq p^2$ for $u,v\in V$. Iterating this and using independence shows

$$\mathbb{P}[T>2t]\leq \mathbb{P}[X_{2t}\neq Y_{2t} \mid X_0=u,Y_0=v]\leq (1-p^2)^t\leq e^{-1}.$$

For some large enough constant t. From Theorem A.1 ([14]) and using the definition of the mixing time we derive $t_{\text{mix}}=O(1)$. We now consider the query complexity w.r.t. $\mathcal{O}^\pi$. By Lemma 4.1 we have that $\left\|\pi_G-\pi_{\widetilde{G}}\right\|_{\mathsf{TV}}=O\left(\left(\frac{k}{n}\right)^{k-1}\right)$. This implies that the first $\Omega\left(\left(\frac{n}{k}\right)^{k-1}\right)$ samples can be coupled with constant probability. Thus, in order to distinguish between G on n nodes and $\widetilde{G}$ on $2n$ nodes one requires $\Omega\left(\left(\frac{n}{k}\right)^{k-1}\right)$ samples. We now consider the query complexity w.r.t. $\overrightarrow{\mathcal{O}}$. In response to the init query, both oracles $\overrightarrow{\mathcal{O}}(G)$ and $\overrightarrow{\mathcal{O}}(\widetilde{G})$ pick v_1, set $\ell(v_1)=1$ and return $(\ell(v_1),\deg^+(v_1))$. Any algorithm has to reach at least node v_k in order to distinguish between G and $\widetilde{G}$. We assume that the oracles $\overrightarrow{\mathcal{O}}(G)$ and $\overrightarrow{\mathcal{O}}(\widetilde{G})$ use independent random adjacency lists at all nodes. But we couple them to be the same for the nodes that are common to both graphs. Clearly, unless $\Omega(n/k)$ queries are made to the out-neighbour lists of each of $v_1,v_2,\ldots,v_{k-1}$, the node v_k cannot be discovered, which in turn means that it is impossible to distinguish between the oracles $\overrightarrow{\mathcal{O}}(G)$ and $\overrightarrow{\mathcal{O}}(\widetilde{G})$. Thus, $\Omega(k\cdot n/k)=\Omega(n)$ queries are required. This concludes the proof. □

Note that the above results only apply to the oracles $\overrightarrow{\mathcal{O}}$ and $\mathcal{O}^\pi$, but not to $\overrightarrow{\mathcal{O}}(1)$, since the in-degrees make it easy to distinguish between the two graphs. However, it is straightforward to extend the graph such that the sample complexity remains $\Omega(n)$ even if the in-degrees are known; thus even with access to $\overrightarrow{\mathcal{O}}(1)$, $\Omega(n)$ queries are required.

Observation 4.3. *Let n be a multiple of k. Then DoubleComet$(2n,2k)$ (Defined in Figure 5) has mixing time $t_{\text{mix}}=O(1)$ and diameter $2k$. Furthermore, any algorithm requires at least $\Omega(n)$ queries to $\overrightarrow{\mathcal{O}}(1)$ to distinguish between $G=$ Comet(n,k) and $\widetilde{G}=$ DoubleComet$(2n,2k)$ on $2n$ nodes.*

The proof is along the same lines as the proof of Theorem 4.2. Observe that the in-degrees of nodes $v_1,\ldots,v_k$ give no extra information about the size of the graph.

4.1 Assuming a Bound on the Connectivity

In this section we introduce the parameter *general conductance*. We first recall some graph notation in the directed setting. Given a non-empty proper subset of vertices $S\subset V$, let $\deg^+(S)=|\{(u,v)\in E: u\in S\}|$ be the out-degree of S. The *cut* of S, ∂S, is the set of edges crossing between S and $V\setminus S$, that is, $\partial S=\{(u,v)\in E: u\in S, v\notin S\}$. The *general conductance* of S, $\phi(S)$, is the ratio between the cut of S, and the out-degree of S. That is, $\phi(S)=|\partial S|/\deg^+(S)$. Given $\varepsilon>0$, the graph ε-general conductance, ϕ_ε, is the minimum of $\phi(S)$ over every non-empty proper subset of V of size at most $(1-\varepsilon)|V|$, *i.e.,*

$\phi_\varepsilon(G) = \min_{S\subseteq V:1\le|S|\le(1-\varepsilon)|V|}\phi(S)$. Note that the parameter ϕ_ε decreases monotonically as ε decreases. In the undirected setting for $\varepsilon = 1/2$ this is just what is commonly known as the conductance.[6]

In the following we describe the algorithm that estimates the graph size.

Upper bound in terms of the general conductance. We consider algorithm EdgeSampling for estimating the number of nodes. The algorithm takes as input the parameter ϕ, a lower bound on the general conductance ϕ_ε. The query complexity is $O(n/\phi_\varepsilon)$ and the output estimate $\widehat{n}$ satisfies $(1-\varepsilon)n \le \widehat{n} \le n$ with arbitrary confidence controlled by an input parameter ℓ. Observe that $O(n/\phi_\varepsilon)$ can be much smaller than the run time of breadth-first search $\Omega(m)$.

Algorithm overview. The algorithm works as follows. At each time step the algorithm maintains a counter Y. If at some point the counter exceeds the threshold ℓ, then the algorithm terminates. The algorithm divides the queries into blocks of length at most $2/\phi$ corresponding to the execution of the **for** loop. In each block, at every step the algorithm samples one outgoing edge uniformly at random from those available and not queried before. If at any step a new node is disclosed, then this finishes the block (break of the **for** loop) and the counter Y is decreased by 1. If the the block finishes without finding a new node, then the counter is increased by 1. Once the counter reaches ℓ, which will happen eventually, then the algorithm outputs the number of nodes it discovered. Even though our goal is to minimise the query complexity, it is worth noticing that the time and space complexity can be kept low by choice of suitable data structures. Although the oracle returns labels of nodes, we use nodes and their labels interchangeably in the algorithm and the analysis of Theorem 4.4.

Theorem 4.4. *Algorithm* EdgeSampling$(\overrightarrow{\mathcal{O}},\ell,\phi)$ *on graph* G *has a query complexity of* $\min\{2(2n+\ell)/\phi, m\}$ *and outputs an estimate* $\widehat{n} \le n$. *Furthermore, if* G *has general conductance* $\phi_\varepsilon(G)$ *of at least* ϕ, *then the algorithm satisfies* $\widehat{n} \ge (1-\varepsilon)n$ *w.p. at least* $1-2^{-\ell}$.

Proof. First observe, that per iteration of the **for** loop at most one new node is discovered. Let Y_t be the variables defined in the algorithm and define $X_t = |S_t| - |S_{t-1}|$. In particular $X_t \in \{0,1\}$. Note that the number of times the counter Y decreases, *i.e.,* $Y_t < Y_{t-1}$, during the first τ iterations of the **for** loop is $\sum_{i\le\tau} X_i \le n$, since the counter increases when a new node is discovered. Therefore, the number of queries is bounded by $2\cdot(2n+\ell)/\phi$. Moreover, every edge is only queried at most once yielding the claimed bound on the query complexity. We now prove that the output $\widehat{n} := |S_t|$ satisfies $\widehat{n} \ge (1-\varepsilon)n$ under the conditions in the statement.

In the remainder we assume $\phi_\varepsilon(G) \ge \phi$ since otherwise the statement is trivially true. Let

$$T = \min\left\{ t \ge 0 \;\middle|\; \sum_{i=0}^{t} X_i \ge n(1-\varepsilon) \text{ or } Y_t = \ell \right\}.$$

Note that every iteration $j \le T$ of the **for** loop satisfies that the probability of finding a new node is at least $|\partial S_j|/\deg^+(S_j) \ge \phi_\varepsilon(G) \ge \phi$. Hence, we have that $\mathbb{P}[X_j = 0] \le (1-\phi)^{2/\phi} < 1/3$. Observe that $X_t = 1$ implies $Y_t = Y_{t-1} - 1$ and $X_t = 0$ implies $Y_t = Y_{t-1} + 1$. Consider the Markov chain $(Z_t)_{t\ge0}$ defined in Proposition A.2 with $s = n, b = n+\ell$ and $p = 1/3$. We couple Y_t and Z_t for $t \le T$ such that $Y_t + n \le Z_t$,

[6] The definition of conductance in the directed setting is more involved and more importantly doesn't seem be directly relevant to the question of estimating the number of nodes. It suffers from similar problems as the skewed stationary distributions. Graphs having poor connectivity to a large fraction of the nodes may still have very *good* conductance if the total mass of the poorly connected nodes under the stationary distribution is very small.

Algorithm 4.1 EdgeSampling($\overrightarrow{\mathcal{O}},\ell,\phi$)

$Y_0=0$ (fail surplus counter)
$v=(node)$ $\overrightarrow{\mathcal{O}}$.init (query oracle to get the initial node)
$S_0=\{v\}$
$E_0=\{(v,i) \mid i\leq \deg^+(v)\}$ (set of undisclosed edges)
$t=1$
while $Y_t\leq \ell$ **do**
  for $\tau=1$ **to** $2/\phi$ **do**
    choose (u,i) uniformly at random from $E_{t,\tau-1}$.
    $v=(node)$ $\overrightarrow{\mathcal{O}}.(u,i,\mathsf{out})$
    if $v\notin S_{t-1}$ **then**
      $S_t=S_{t-1}\cup\{v\}$
      $E_{(t-1)\cdot 2/\phi+\tau}\leftarrow (E_{(t-1)\cdot 2/\phi+\tau-1}\cup\{(v,i) \mid i\leq \deg^+(v)\})\setminus\{(u,i)\}$
      break
    else
      $E_{t\cdot 2/\phi}\leftarrow E_{t\cdot 2/\phi-1}\setminus\{(u,i)\}$
  if $|S_t|=|S_{t-1}|+1$ **then**
    $Y_t\leftarrow Y_t-1$
  else
    $Y_t\leftarrow Y_t+1$
    $S_t\leftarrow S_{t-1}$
  $t\leftarrow t+1$
Output $|S_t|$.

since $\mathbb{P}[Y_t=Y_{t-1}+1]=\mathbb{P}[X_t=0]\leq 1/3=p=\mathbb{P}[Z_t=Z_{t-1}+1]$. Let

$$T'=\min\Big\{t\geq 0 \;\Big|\; Z_t\in\{0,n+\ell\}\Big\}.$$

From Proposition A.2 we get that

$$\mathbb{P}[Z_{T'}=n+\ell]=\frac{\left(\frac{2/3}{1/3}\right)^n-1}{\left(\frac{2/3}{1/3}\right)^{n+\ell}-1}\leq 2^{-\ell}.$$

Due to our coupling we have $Y_t+n\leq Z_t$ and observe that $Z_t=0$ implies that $\sum_i^t X_i\geq n$ and therefore we have $T\leq T'$. Hence, $Z_{T'}=0$ implies $Z_T<n+\ell$ which in turn implies $Y_T<\ell$. Thus,

$$\begin{aligned}
\mathbb{P}[|S_T|\geq n(1-\varepsilon)] &= \mathbb{P}\left[\sum_{i=0}^{T} X_i\geq n(1-\varepsilon)\right]\\
&= \mathbb{P}[Y_T<\ell]\\
&\geq \mathbb{P}[Z_T<n+\ell]\\
&\geq \mathbb{P}[Z_{T'}<n+\ell]\\
&= 1-\mathbb{P}[Z_{T'}=n+\ell]\\
&\geq 1-2^{-\ell},
\end{aligned}$$

which concludes the proof. □

Observe that the error made by Algorithm EdgeSampling is one-sided — the estimate never exceeds n. Allowing a two-sided error and given knowledge of ε, one can instead output an estimate, which has a smaller additive error. This is summarised in the following observation.

Observation 4.5. *Consider a modification of Algorithm* EdgeSampling$(\vec{\mathcal{O}},\ell,\phi)$ *on graph G which takes the additional parameter ε and outputs $\widehat{n}^* := |S_t|(1+\frac{\varepsilon}{2-\varepsilon})$ instead of $|S_t|$. If G has general conductance $\phi_\varepsilon(G)$ of at least ϕ, then $\widehat{n}^*$ satisfies $|n-\widehat{n}^*| \leq \frac{\varepsilon}{2-\varepsilon}n$ w.p. at least $1-2^{-\ell}$.*

Proof. By Theorem 4.4 we have $n \geq |S_t| \geq (1-\varepsilon)n$, where t is the the index of the last execution of the **while** loop. We have that

$$\begin{aligned} |n-\widehat{n}^*| &= \left|n-|S_t|\left(1+\frac{\varepsilon}{2-\varepsilon}\right)\right| \\ &\leq \max\left\{\left|n-n\left(1+\frac{\varepsilon}{2-\varepsilon}\right)\right|, \left|n-n(1-\varepsilon)\left(1+\frac{\varepsilon}{2-\varepsilon}\right)\right|\right\} \\ &= \frac{\varepsilon}{2-\varepsilon}n, \end{aligned}$$

which finishes the proof. □

Lower bound in terms of the general conductance. In the following we show that the bound of Observation 4.5 is almost tight. Recall that, given ϕ_ε, the modified version of Algorithm EdgeSampling in Observation 4.5 returns an estimate with and additive error of at most $\frac{\varepsilon}{2-\varepsilon}n$ using $O(n/\phi_\varepsilon)$ queries. In what follows we show that any algorithm, given the values ϕ_ε and ε, cannot output an estimate with an error smaller than $\frac{\varepsilon-\delta}{2-\varepsilon-\delta}n$ unless it makes $\Omega(n/\phi_\varepsilon)$ queries, for any $\delta < \varepsilon/2$. We prove the following lemma for undirected graphs using $\mathcal{O}$, but it should be clear that the same the same result holds for directed graphs, by making the graph directed, with symmetric edges, and using $\vec{\mathcal{O}}(2)$ (and hence also for oracles $\vec{\mathcal{O}}$ and $\vec{\mathcal{O}}(1)$).

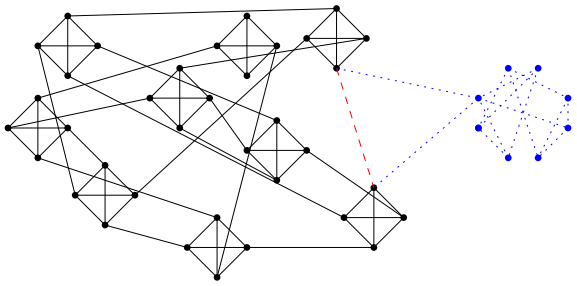

Figure 6: The graphs of Theorem 4.6. G contains the black nodes and the black and red (dashed) edges which form a $\lceil 1/\phi \rceil$-regular expander on cliques of size $\lceil 1/\phi \rceil$. G' is obtained by removing the red (dashed) edge and adding the blue (dotted) graph. At least one blue edge needs to be sampled, which takes $\Omega(n/\phi)$ time, in order to estimate n accurately.

Theorem 4.6. *Let $n \in \mathbb{N}$, $\phi \in [1/n, 1]$ and $\varepsilon \in (0, 1/2]$. There exists an undirected graph with general conductance $\phi_\varepsilon = \Theta(\phi)$ such that any algorithm with access to $\mathcal{O}$ requires $\Omega(n/\phi_\varepsilon)$ queries to output $\widehat{n}$ such that $|n-\widehat{n}| \leq \frac{\varepsilon-\delta}{2-\varepsilon-\delta}n$ w.p. at least $2/3$ for any $\delta < \varepsilon/2$.*

Proof. Let $d=\lceil 1/\phi \rceil$. For simplicity we assume that d divides n. Let $G_1=(V_1,E_1)$ be a d-regular expander on $n^*=n(1-\varepsilon+\delta/2)/d$ nodes. We construct $G=(V,E)$ (see Figure 6 for an illustration) by replacing every node r by a clique K of size d such that every node of the clique has exactly one of the edges incident to r. Let $n_1:=|V|=n(1-\varepsilon+\delta/2)$. Let G_2 be a 3-regular expander on $n-n_1$ nodes. Choose an edge (u,v) uniformly at random among all edges where both nodes u and v belong to the same clique. We construct G' by removing (u,v), and by adding the edge (u,w),(v,w), where w is an arbitrary node of G_2.

The proof idea is to couple the oracle decisions such that w.p. $2/3$ no algorithm can distinguish between G and G'. Incidentally, our construction ensures that there is no estimate $\widehat{n}$ satisfying $|\widehat{n}-n_1| \leq \frac{\varepsilon-\delta}{2-\varepsilon-\delta} n_1$ and $|\widehat{n}-n| \leq \frac{\varepsilon-\delta}{2-\varepsilon-\delta} n$ simultaneously. Suppose for the sake of contradiction this was possible, then $\widehat{n}-n(1-\varepsilon+\delta/2) \leq \frac{\varepsilon-\delta}{2-\varepsilon-\delta} n_1$ and $n-\widehat{n} \leq \frac{\varepsilon-\delta}{2-\varepsilon-\delta}$. Adding up both equations yields $n-n(1-\varepsilon+\delta/2) \leq \frac{\varepsilon-\delta}{2-\varepsilon-\delta}(n_1+n)$. This implies $\varepsilon-\delta/2 \leq \frac{\varepsilon-\delta}{2-\varepsilon-\delta}(2-\varepsilon+\delta/2)$, which can be shown to a contradiction for $\varepsilon \leq 1/2$. Thus, any algorithm approximating n up to an error of $\frac{\varepsilon-\delta}{2-\varepsilon-\delta}$ needs to distinguish between G and G'.

We assume that the oracles $\mathcal{O}(G)$ and $\mathcal{O}(G')$ use a consecutive labelling function, *i.e.,* whenever it chooses a new node, it picks the smallest as yet unused natural number as the label. In response to the init query, both oracles pick $v_1 \in V_1 \setminus \{u,v\}$, set $\ell(v_1)=1$ and return $(\ell(v_1),\deg_G(v_1))$. As long as the algorithm doesn't disclose the edge (u,w) or (v,w) in G' and (u,v) in G, we can couple the nodes returned by the oracle and thus the two graphs are indistinguishable. Hence, similarly as before, we couple with constant probability for the first cn/ϕ samples returned by the oracle for constant $c>0$ small enough. It remains to show that G' has general conductance $\phi_\varepsilon=\Theta(\phi)$.

Consider an arbitrary clique K in the construction. Clearly, $\phi_\varepsilon \leq \phi(V(K))=O(d/d^2)=O(\phi)$. To derive a lower bound on ϕ_ε consider any set S such that $\phi=\phi(S)$. Suppose that S contains only entire cliques, *i.e.,* if $u \in S$, then also $v \in S$ for u,v belonging to the same clique. Then, the general conductance of S is $\phi(S)=\Theta(\phi^{G_1}/d)=\Theta(\phi)$, where $1/\phi^{G_1}=\Theta(1)$ is the general conductance of the G_1.

We claim that there exists a set of cliques, $\mathcal{K}$, containing a constant fraction of the cliques and no node of S is in any clique of $\mathcal{K}$. If this were not the case, then S either violates $|S| \leq n(1-\varepsilon)$ or S violates $\phi=\phi(S)$.

Now suppose that S does not contain only entire cliques. Observe that for any clique K which is not entirely included in S, we can include the remaining nodes which only decreases the general conductance. The new resulting set S' might be larger than $n(1-\varepsilon)$. However, observe that $\phi_\varepsilon \geq \phi_{\varepsilon-\varepsilon'}$ for any sufficiently small $\delta > \varepsilon' > 0$. Note that S' does not contain any nodes of cliques in $\mathcal{K}$. Hence, the resulting general conductance fulfils $\phi(S) \geq \phi(S') \geq \Omega(\phi^{G_1}/d)=\Omega(\phi)$. This completes the proof. □

References

[1] Michael A. Bender and Data Ron. Testing properties of directed graphs: Acyclicity and connectivity. *Random Structures and Algorithms*, 20(2):184–205, 2002.

[2] Béla Bollobás. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European Journal of Combinatorics*, 1(4):311 – 316, 1980.

[3] Colin Cooper and Alan Frieze. Crawling on web graphs. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 419–427. ACM, 2002.

[4] Colin Cooper, Tomasz Radzik, and Yiannis Siantos. Estimating network parameters using random walks. In *Computational Aspects of Social Networks, 4th International Conference on*, pages 33–40, 2012.

[5] Artur Czumaj, Pan Peng, and Christian Sohler. Relating two property testing models for bounded degree directed graphs. In *Proceedings of the ACM Symposium on the Theory of Computing (STOC)*, 2016.

[6] Anirban Dasgupta, Ravi Kumar, and Tamás Sarlós. On estimating the average degree. In *Proceedings of the 23rd international conference on World Wide Web*, pages 795–806. ACM, 2014.

[7] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, 1968.

[8] Mark Finkelstein, Howard G Tucker, and Jerry Alan Veeh. Confidence intervals for the number of unseen types. *Statistics & Probability Letters*, 37(4):423–430, 1998.

[9] Oded Goldreich. Introduction to testing graph properties. Survey article available at http://www.wisdom.weizmann.ac.il/~oded/COL/tgp-intro.pdf, 2010.

[10] I. J. Good. The population frequencies of species and the estimation of population parameters. *Biometrika*, 40(3/4):237–264, 1953.

[11] Liran Katzir and Stephen J. Hardiman. Estimating clustering coefficients and size of social networks via random walk. *ACM Transactions on the Web*, 9(4):19:1–19:20, 2015.

[12] Liran Katzir, Edo Liberty, Oren Somekh, and Ioana A. Cosma. Estimating sizes of social networks via biased sampling. *Internet Mathematics*, 10(3-4):335–359, 2014.

[13] Donald E. Knuth. Estimating the efficiency of backtrack programs. *Mathematics of computation*, 29(129):122–136, 1975.

[14] David Levin, Yuval Peres, and Elizabeth Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2006.

[15] Alberto Marchetti-Spaccamela. On the estimate of the size of a directed graph. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 317–326. Springer, 1988.

[16] Milena Mihail, Amin Saberi, and Prasad Tetali. Random walks with lookahead on power law random graphs. *Internet Mathematics*, 3(2):147–152, 2006.

[17] Cameron Musco, Hsin-Hao Su, and Nancy A. Lynch. Ant-inspired density estimation via random walks: Extended abstract. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*, pages 469–478, 2016.

[18] Gregory Valiant and Paul Valiant. Estimating the unseen: an n/log(n)-sample estimator for entropy and support size, shown optimal via new clts. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 685–694. ACM, 2011.

[19] Gregory Valiant and Paul Valiant. Estimating the unseen: improved estimators for entropy and other properties. In *Advances in Neural Information Processing Systems*, pages 2157–2165, 2013.

A Auxiliary Claims

Theorem A.1 ([14]). *Consider two irreducible Markov chains $(X_t)_{t\geq 0}$, $(Y_t)_{t\geq 0}$ with transition matrix P, $X_0=x$, and $Y_0=y$. Let $\{(X_t,Y_t)\}$ be a coupling satisfying that if $X_s=Y_s$, then $X_t=Y_t$ for $t\geq s$. Let $T=\min\{t\geq 0 \mid X_t=Y_t\}$. Then $\left\|p^t(x,\cdot)-p^t(y,\cdot)\right\|_{\mathsf{TV}}\leq \mathbb{P}[T>t]$.*

Proposition A.2. *[7, Chapter XIV.2]] Let $p\in(0,1/2)$ and $b,s\in\mathbb{N}$. Consider a discrete time Markov chain $(Z_t)_{t\geq 0}$ with state space $\Omega=[0,b]$ where*

- $Z_0=s\in[0,b]$
- $\mathbb{P}[Z_t=i \mid Z_{t-1}=i-1]=p$ *for* $i\in[1,b-1],t\geq 1$
- $\mathbb{P}[Z_t=i \mid Z_{t-1}=i+1]=1-p$ *for* $i\in[1,b-1],t\geq 1$
- $\mathbb{P}[Z_t=i \mid Z_{t-1}=i]=1$ *for* $i\in\{0,b\},t\geq 1$

Let $T=\min\{t\geq 0 \mid Z_t\in\{0,b\}\}$. Then

$$\mathbb{P}[Z_T=b]=\frac{\left(\frac{1-p}{p}\right)^s-1}{\left(\frac{1-p}{p}\right)^b-1}.$$